RESUMING COMMUNICATION OF CONTENT DATA ON THE INITIAL CHANNEL WHEN THE INITIAL AND THE NEW ACCESS POINT PERFORM HANDOVER STEPS

Field of the Invention

The present invention relates to wireless access networks, which include wireless local area networks (WLANs) and wireless wide area networks (WWANs), and relates in particular to handovers in wireless access networks. The present invention relates in particular, but not exclusively, to handover of a mobile node (MN) between two access points (APs) under the Inter Access Point Protocol (IAPP) in a WLAN operating according to the IEEE 802.11 standard.

Background of the Invention

Local area networks (LANs) are computer and communications networks which users may access at different access points (APs). One type of LAN is a wireless local area network (WLAN) in which user's equipment may be mobile and connects with an access point by means of a wireless link. A well known operating standard for WLANs is IEEE 802.11.

Mobile user equipment in a WLAN is usually referred to as a mobile node (MN). Some examples of an MN are a mobile telephone with a WLAN interface, a laptop computer with a WLAN interface, and a personal data assistant (PDA) with a WLAN interface.

When a MN moves, a requirement can arise for the MN to be handed over from one AP (which may be called the "initial" or "old" AP) to a "new" AP. Handover is controlled by a handover protocol. In the case of a WLAN operating according to IEEE 802.11, the handover protocol is typically the Inter Access Point Protocol (IAPP), which applies across the "initial" and the "new" AP.

The IAPP provides a means for transferring so-called "context" from the old AP to the new AP. The term context is used to describe information relating to the MN and its operation that is required at an AP for the AP to provide service to the MN. Such information may include security information such as authentication information, quality of service (QoS) information, and so on. However, in typical WLANs, and operating standards thereof, handover causes a break in service before the new AP takes over serving the MN. This may be disadvantageous, particularly when the MN is running a real-time application, e.g. a real-time multi-media application such as receiving and displaying real-time video.

In some WLANs, and operating standards thereof, so-called soft handover processes (in which two or more APs effectively serve the MN at the same time) have been applied in an attempt to alleviate the disadvantages of a break in service due to handover.

However, in many WLANs, or operating standards thereof, soft handover is not possible and/or not desirable. This is the case for WLANs operating under

the IEEE 802.11 standard, in which adjacent APs typically operate in different frequency bands.

Also, a problem related to QoS provision may arise with handover of real-time applications in WLANs. In particular, where details of operating standards are based upon standards designed originally for fixed networks, a process of requesting and obtaining a required QoS level, using so-called reservation paths, tends to be problematic in WLANs where a user is allowed to change access points, essentially due to the need for re-establishing reservation paths after handover. For example, an Integrated Services (IntServ) framework, as standardized by the Internet Engineering Task Force (IETF) and discussed in R. Braden et al., "Integrated Services in the Internet Architecture: an Overview," RFC 1633, June 1994, provides a means for requesting and obtaining QoS per flow. IntServ uses Resource ReserVation Protocol (RSVP), as described in R. Braden et al., "Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification," RFC2205, September 1997, for implementing signalling associated with this. However, having been designed originally for fixed networks, RSVP is an example of a process which tends to be problematic in WLANs. In particular, depending on the distance between the peers, considerable delays can arise, deteriorating the network performance during handovers, especially for real-time applications.

Considerations such as those discussed above with respect to WLANs also apply to wireless wide area networks (WWANs). WLANs and WWANs may be considered to be two types of wireless access networks, and in this specification

the term "wireless access network" is to be understood to include at least WLANs and WWANs, as well as any other access networks whose characteristics are able to benefit from the following invention.

Summary of the Invention

In a first aspect, the present invention provides a method of changing access points for a mobile node, or performing handover, in a wireless access network, for example a wireless local area network or a wireless wide area network, the method comprising: the mobile node communicating content data on an initial channel via an initial access point of the wireless access network; the mobile node sending a handover request on a new channel to a new access point of the wireless access network; the mobile node resuming communication of content data on the initial channel via the initial access point; elements of the wireless access network, including the initial access point and the new access point, performing handover steps while the mobile node performs the resumed communication of content data on the initial channel via the initial access point; and when the handover steps are completed, the mobile node communicating content data on the new channel via the new access point.

The step of the mobile node performing the resumed communication of content data on the initial channel via the initial access point may end by virtue of the mobile node switching from the initial channel to the new channel in response to an instruction message from the wireless access network which is

sent as part of the handover steps performed by elements of the wireless access network.

The method may further comprise the wireless access network sending a completion message to the mobile node informing the mobile node that the handover steps are completed.

The step of the mobile node communicating content data on the new channel via the new access point may start in response to the mobile node receiving the completion message.

The step of the mobile node switching from the initial channel to the new channel in response to a message from the wireless access network which is sent as part of the handover steps performed by elements of the wireless access network may be performed a predetermined or a calculated amount of time after the mobile node receives the instruction message.

The method may further comprise the step of the mobile node determining whether an expected data flow of the new access point is acceptable, for example in terms of quality of service.

In a further aspect, the present invention provides a storage medium storing processor-implementable instructions for controlling a processor to carry out the method steps of any of the above aspects performed by the mobile node.

In a further aspect, the present invention provides a storage medium storing processor-implementable instructions for controlling a processor to carry out the method steps of any of the above aspects performed by the access points.

In a further aspect, the present invention provides a mobile node adapted to perform the method steps of any of the above aspects performed by the mobile node.

In a further aspect, the present invention provides an access point adapted to perform the method steps of any of the above aspects performed by the access points.

In a further aspect, the present invention provides apparatus for changing access points for a mobile node, or performing handover, in a wireless access network, for example a wireless local area network or a wireless wide area network, comprising: means for the mobile node to communicate content data on an initial channel via an initial access point of the wireless access network; means for the mobile node to send a handover request on a new channel to a new access point of the wireless access network; means for the mobile node to resume communication of content data on the initial channel via the initial access point; means for elements of the wireless access network, including the initial access point and the new access point, to perform handover steps while the mobile node performs the resumed communication of content data on the initial channel via the initial access point; andmeans for, when the handover steps are completed,

the mobile node to communicate content data on the new channel via the new access point.

The means for the mobile node to perform the resumed communication of content data on the initial channel via the initial access point may be arranged to end the resumed communication by virtue of the mobile node switching from the initial channel to the new channel in response to an instruction message from the wireless access network which is sent as part of the handover steps performed by elements of the wireless access network.

The apparatus may further comprise means for the wireless access network to send a completion message to the mobile node informing the mobile node that the handover steps are completed.

The means for the mobile node to communicate content data on the new channel via the new access point may be arranged to start communication of the content data in response to the mobile node receiving the completion message.

The means for the mobile node to switch from the initial channel to the new channel in response to a message from the wireless access network which is sent as part of the handover steps performed by elements of the wireless access network may be arranged such that the switching is performed a predetermined or a calculated amount of time after the mobile node receives the instruction message.

The apparatus may further comprise means for the mobile node to determine whether an expected data flow of the new access point is acceptable.

In any of the above aspects, the wireless access network may be a wireless local area network operating according to the operating standard IEEE 802.11.

The present invention tends to alleviate or resolve problems associated with a break occurring between service from an old AP and a new AP, by the MN sending a request to a new AP but then reverting to the old AP until handover preparation in a fixed part of the WLAN is complete, or at least substantially complete, and the MN is informed of this. This tends to reduce handover latency between the old AP and the new AP and avoids breaks in service, thus improving in particular provision of real-time applications such as real-time multi-media applications.

Brief Description of the Drawings

Embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

FIG. 1 is a schematic illustration of a wireless local area network (WLAN);

FIG. 2 is a call sequence flowchart showing process steps employed in a prior art handover process; and

FIG. 3 is a call sequence flowchart showing process steps employed in a handover process according to an embodiment of the invention.

Description of Preferred Embodiments

FIG. 1 is a schematic illustration of a wireless local area network (WLAN) 1. The WLAN 1 comprises a transport network, which in this example is an IP network 2, coupled to an authentication unit 3 and a plurality of access points (AP's). The access points are coupled, via radio links, to various mobile nodes (MN's). In practise the WLAN 1 will typically comprise a large number of AP's and MN's, but for clarity only two AP's, namely a first AP 4 and a second AP 5 are shown, and only one MN, namely a MN 6, are shown. When the MN 6 is coupled to the first AP 4 this is by a radio link 24; when the MN 6 is coupled to the second AP 5 this is by a radio link 25. In this example the MN 6 is a laptop computer with a WLAN interface, but in general may be any known type, for example a mobile telephone with a WLAN interface or a personal data assistant (PDA) with a WLAN interface.

In this embodiment the WLAN 1 operates according to the well known operating standard IEEE 802.11. The authentication unit 3 is a Remote Authentication Dial-In User Service (RADIUS) server.

The WLAN 1 allows the MN 6 to communicate with an entity via the IP network 2. Such entity may be another user node of the WLAN coupled to the IP network 2 via an AP. Another possibility is that the IP network 2 may be coupled

to a gateway to a public switched telephone network (PSTN), and the entity is connected via the PSTN. The type and connection route of the entity is unimportant to understanding this embodiment, and therefore in FIG. 1 and the following description this connection is presented in general terms as a data route 8 coupled to the IP network 2.

In operation, the MN 6 interacts with the IP network 2 via an access point. More particularly, let us consider a scenario where the MN 6 is engaged in a data session, and access is initially provided via, say, AP 4, which will therefore be termed hereinafter the initial AP 4. The data required to pass between various elements may be divided into data representing the information content of the data session (hereinafter referred to as content data), and signalling data used to set up and maintain the flow of the content data between elements. Content data passes between the MN 6 and the initial AP 4, between the initial AP 4 and the IP network 2, and between the IP network 2 and the data route 8. Handover specific signalling data, including encryption data, passes between the MN 6 and the initial AP 4, between the initial AP 4 and the IP network 2, and between the IP network 2 and the authentication unit 3.

Let us now consider when it is required or desired for the MN 6's access to the communications network 2 to be changed so that it then becomes provided by a different AP, e.g. the AP 5, which will therefore be termed hereinafter the new AP 5. The implementation of this is called handover. Conventionally, this is performed using a handover mechanism, comprising for example the Inter Access Point Protocol (IAPP) and the IEEE 802.11 standard.

In this embodiment, the AP's 4, 5 and the MN 6 have been adapted, to offer, and provide for, an adapted form of access handover, as will be described in more detail below.

The adaptation may be implemented in the respective elements in any suitable manner. For example, new apparatus may be added to conventional AP's and MN's, or alternatively existing parts of a conventional AP's and MN's may be adapted, for example by reprogramming of a one or more processors therein. As such the required adaptation may be implemented in the form of processor-implementable instructions stored on a storage medium, such as a floppy disk, hard disk, PROM, RAM or any combination of these or other storage media.

The adapted form of handover of this embodiment may be most readily understood by comparison with the conventional handover process, which is therefore described first, as follows.

FIG. 2 is a call sequence flowchart showing process steps employed in the prior art handover process. In the case of the MN 6, two separate radio channels are shown, an initial channel 14 which is the channel used by the MN 6 to communicate with the initial AP 4, and a new channel 15 which is the channel used by the MN 6 to communicate with the new AP 5. The channels are different radio frequencies from each other, as specified in IEEE Std. 802.11.

Initially, a data session is underway, via initial AP 4. This is represented in FIG. 2 by step 2 in which content data is passed between the MN6 operating on the initial channel 14, the initial AP 4, the IP network 2 and the data route 8. In more detail, step 2 comprises: a step 2a in which content data is passed between the MN 6 operating on its initial channel 14 and the initial AP 4; a step 2b in which content data is passed between the initial AP 4 and the IP network 2; and a step 2c in which content data is passed between the IP network 2 and the data route 8. Although presented as a discrete step 2 for consistency in FIG. 2, it will be understood that this represents an ongoing two-way data flow process. This data flow process is shown in FIG. 2 as being undertaken during a first time period 21. The active channel of the MN 6 during the first time period 21 is the initial channel 14.

Then handover is performed by virtue of steps s4 to s18 as will be described below. Steps s4 to s18 are implemented over the course of a second time period 22, as shown in FIG. 2. The active channel of the MN 6 during the second time period 22 is the new channel 15. Thus, during the second time period 22 the flow of content data is interrupted (i.e. step 2 does not take place during the second time period), since the MN 6 has switched to its new channel 15, which corresponds to the new AP 5 rather than the initial AP 4, but the MN 6 is not yet able to send content data to or receive content data from the new AP 5.

Steps s4 to s18 follow the protocol specified in the IEEE Std. 802.11 and the Inter Access Point Protocol (IAPP), which is a recommended practice currently standardized by the IEEE Working Group 802.11f, and is specified in IEEE Std.

802.11f/Draft 5, "Draft Recommended Practice for Multi-Vendor Access Point Interoperability via an Inter-Access Point Protocol Across Distribution Systems Supporting IEEE 802.11 Operation". The IAPP provides capabilities for transferring context from one AP to another, through the fixed distribution system, in order to facilitate handovers across multi-vendor APs. (The term context is used to describe information relating to the MN and its operation that is required at an AP for the AP to provide service to the MN.) The IAPP describes a service access point (SAP), service primitives, a set of functions and a protocol that will allow conformant APs to interoperate on a common distribution system, using the Transmission Control Protocol over IP (TCP/IP, where IP is Internet Protocol) to exchange IAPP packets. Both IP in general, and more particularly TCP/IP are well known to the person skilled in the art, hence will only be briefly summarised here to the extent that, broadly speaking, they are protocols that define procedures for exchange of packetized information, and are specified in terms of different hierarchical layers, e.g. Layer-2, layer 3, according to different types of entities or functions needing to be consistent with each other.

Referring to FIG. 2, at step s4, the message exchange during handover is initiated by the MN 6 sending a Reassociate-Request message from its new channel 15 to the new AP 5. The Reassociate-Request message carries the Basic Service Set Identifier (BSSID) of the initial AP 4.

The new AP 5 will need to communicate with the initial AP 4. In order to do this, the new AP 5 will need the IP address of the initial AP 4. In general, mapping between BSSIDs and IP addresses can be either stored at the AP's

themselves, or at an authentication unit such as authentication unit 3. In this example, the mapping is held at the authentication unit. Therefore, at step s6, the new AP 5 sends an Access-Request message to the authentication unit 3, this message including the information of the BSSID of the initial AP 4.

The authentication unit 3 processes this request, including looking up the IP address of the initial AP 4 that it has mapped for the BSSID of the initial AP 4. At step s8, the authentication unit sends an Access-Accept message to the new AP 5. The Access-Accept message includes the IP address of the initial AP 4.

One option is for the communication between the AP's to be encrypted or otherwise made secure. This is the case in this example. Hence, the Access-Accept message sent, at step s8, from the authentication unit to the new AP 5 also includes a Security Key to enable the new AP 5 to prepare a security block for sending to the initial AP 4.

Then security blocks are exchanged between the new AP 5 and the initial AP 4. This is implemented as follows. At step s10, the new AP 5 sends a Security Block to the initial AP 4. At step s12, the initial AP 4 sends a Security Block Acknowledgment to the new AP 5.

At step s14, the new AP 5 sends a Move-notify message to the initial AP 4. This Move-notify message in effect informs the initial AP 4 that the MN 6 has requested to handover to the new AP 5.

At step s16, the initial AP 4 sends a Move-response message to the new AP 5. This Move-response message and the Move-notify message each include a Context Block. Both the Move-notify message sent at step s14 and the Move-response message sent at step s16 are IP packets carried in a TCP session between the two APs. These messages may contain authentication information that allows the new AP 5 to accept the MN 6 without re-authentication. However, the Context Block has a flexible structure which is able to support a range of information exchange. More specifically, it can consist of a variable number of Information Elements (IEs) of the form (Element ID, Length, Information). In this way, every IE can contain variable length information, whose type is specified by the Element ID. Processing of the information transferred inside the IEs is beyond the scope of the IAPP, as it depends on the functionality of the APs, and will be arranged according to the requirements or circumstances of a particular system or set up, as specified by the skilled person.

At step s18, the new AP 5 sends a Layer-2 (L2) Update Frame to the IP network 2. In response to this, the IP network 2 changes the data path to being switched via the new AP 5 instead of the initial AP 4. This is represented in FIG. 2 by a step s20 in which the content data is passed between the MN6 operating on the new channel 15, the new AP 5, the IP network 2 and the data route 8. In more detail, step s20 comprises: a step s20a in which content data is passed between the MN 6 operating on its new channel 15 and the new AP 5; a step s20b in which content data is passed between the new AP 5 and the IP network 2; and a step s20c in which content data is passed between the IP network 2 and the data route 8. Although presented as a discrete step s20 for consistency in FIG. 2,

it will be understood that this represents an ongoing two-way data flow process during a third time period 23, as shown in FIG. 20.

The adapted form of handover of this embodiment will now be described with reference to FIG. 3.

FIG. 3 is a call sequence flowchart showing process steps employed in the handover process of this embodiment. The same reference numerals as were used with respect to FIG. 2 are used again in FIG. 3 to indicate the two separate radio channels, i.e. the initial channel 14 and the new channel 15 used by the MN 6. These channels are as described above with reference to FIG. 2.

In the process of this embodiment, each of the prior art steps described above with reference to FIG. 2 is carried out, but additional steps are now included. For the sake of completeness, the steps common to the prior art approach are renumbered and described again in the following description of this embodiment made with reference to FIG. 3. Individually, these steps corresponding to those present in the above described prior art process again follow the IEEE 802.11 standard and the protocol specified in the Inter Access Point Protocol (IAPP), which is a recommended practice currently standardized by the IEEE Working Group 802.11f, and is specified in IEEE Std. 802.11f/Draft 5, "Draft Recommended Practice for Multi-Vendor Access Point Interoperability via an Inter-Access Point Protocol Across Distribution Systems Supporting IEEE 802.11 Operation".

Initially, a data session is underway, via initial AP 4. This is represented in FIG. 3 by step s30 in which content data is passed between the MN6 operating on the initial channel 14, the initial AP 4, the IP network 2 and the data route 8. In more detail, step s30 comprises: a step s30a in which content data is passed between the MN 6 operating on its initial channel 14 and the initial AP 4; a step s30b in which content data is passed between the initial AP 4 and the IP network 2; and a step s30c in which content data is passed between the IP network 2 and the data route 8. Although presented as a discrete step s30 for consistency in FIG. 3, it will be understood that this represents an ongoing two-way data flow process. This data flow process is shown in FIG. 3 as being undertaken during a first time period 31. The active channel of the MN 6 during the first time period 31 is the initial channel 14.

Then handover is performed by virtue of steps s34, s36, s38, s40, s42, s44, s46, s48, s50, s52, s54 and s56 as will be described below. In this embodiment the implementation of these steps is conveniently considered as being divided over the course of a second time period 32, a third time period 33, and a fourth time period 34, as shown in FIG. 3 and as will be described in more detail below.

Referring to FIG. 3, at step s34, the message exchange during handover is initiated by the MN 6 sending a Reassociate-Request message from its proposed new channel 15 to the new AP 5. The Reassociate-Request message carries the Basic Service Set Identifier (BSSID) of the initial AP 4.

This step s34 is implemented over the course of the second time period 32. The active channel of the MN 6 during this second time period 32 is the new channel 15. Thus, during the second time period 32 the flow of content data is, strictly speaking, interrupted (i.e. step s30 does not take place during the second time period), since the MN 6 has switched to its new channel 15, which corresponds to the new AP 5 rather than the initial AP 4.

However, this will usually only be for a very small time compared to the prior art process, because after implementing step s34 by sending the Reassociate-Request on the proposed new channel 15, the MN 6 reverts to the initial channel 14 to resume transmission and reception of content data on the initial channel 14.

The resumed transmission and reception (flow) of content data on the initial channel 14 takes place over the course of a third time period 33. This takes place on an ongoing basis throughout the third time period 33, but for clarity is schematically represented in FIG. 3 by means of a step s35 shown at the start of the third time period 33 and a step s51 shown shortly before the end of the third time period 33. In both step s35 and step s51, content data is passed between the MN6 operating on the initial channel 14, the initial AP 4, the IP network 2 and the data route 8. In more detail, step s35 comprises: a step s35a in which content data is passed between the MN 6 operating on its initial channel 14 and the initial AP 4; a step s35b in which content data is passed between the initial AP 4 and the IP network 2; and a step s35c in which content data is passed between the IP network 2 and the data route 8. Likewise, step s51 comprises: a step s51a in

which content data is passed between the MN 6 operating on its initial channel 14 and the initial AP 4; a step s51b in which content data is passed between the initial AP 4 and the IP network 2; and a step s51c in which content data is passed between the IP network 2 and the data route 8. The active channel of the MN 6 during the third time period 33 is the initial channel 14.

Whilst the above described data content represented by steps s35 and s51 takes place, the new AP 5, the initial AP 4, the IP network 2 and the authentication unit 3 carry out handover steps as will now be described.

In this embodiment, the mapping is again held at the authentication unit. Therefore, at step s36, the new AP 5 sends an Access-Request message to the authentication unit 3, this message including the information of the BSSID of the initial AP 4.

The authentication unit 3 processes this request, including looking up the IP address of the initial AP 4 that it has mapped for the BSSID of the initial AP 4. At step s38, the authentication unit sends an Access-Accept message to the new AP 5. The Access-Accept message includes the IP address of the initial AP 4.

In this embodiment the option for the communication between the AP's to be encrypted or otherwise made secure is again employed. Hence, the Access-Accept message sent, at step s38, from the authentication unit to the new AP 5 also includes a Security Key to enable the new AP 5 to prepare a security block for sending to the initial AP 4.

Then security blocks are exchanged between the new AP 5 and the initial AP 4. This is implemented as follows. At step s40, the new AP 5 sends a Security Block to the initial AP 4. At step s42, the initial AP 4 sends a Security Block Acknowledgment to the new AP 5. (Note, in other embodiments where encryption is not employed, steps s40 and s42 may be omitted.)

At step s44, the new AP 5 sends a Move-notify message to the initial AP 4. This Move-notify message in effect informs the initial AP 4 that the MN 6 has requested to handover to the new AP 5.

At step s46, the initial AP 4 sends a Move-response message to the new AP 5. This Move-response message and the Move-notify message each include a Context Block. Both the Move-notify message sent at step s44 and the Move-response message sent at step s46 are IP packets carried in a TCP session between the two APs. These messages may contain authentication information that allows the new AP 5 to accept the MN 6 without re-authentication. However, the Context Block has a flexible structure which is able to support a range of information exchange. More specifically, it can consist of a variable number of Information Elements (IEs) of the form (Element ID, Length, Information). In this way, every IE can contain variable length information, whose type is specified by the Element ID. Processing of the information transferred inside the IEs is beyond the scope of the IAPP, as it depends on the functionality of the APs, and will be arranged according to the requirements or circumstances of a particular system or set up, as specified by the skilled person.

The Move-Response message sent at step s46 is based on and includes the contents of this message that are present under conventional IAPP signalling. By virtue of this, the new AP 5 receives context information pertaining to the MN 6, which may include for example security context and Quality of Service (QoS) context. Such QoS context may include both IP layer-2 QoS attributes, which are applicable to the WLAN radio interface, and IP layer-3 QoS attributes, which specify end-to-end QoS requirements. Preferably a suitable layer-2 QoS management scheme is employed, to provide for different kinds of traffic to receive acceptable service in the wireless link.

Additionally, in this embodiment, for the sake of IP layer-3 QoS management, the Move-Response message sent at step s46 further includes an RSVP context. This RSVP context contains information about at least some, but preferably all, of the active IP flows, i.e. the data content and signal data flows being transmitted and received as part of the session. The new AP 5 might already have more current data flow than the initial AP 4, or at least relatively so compared to its capacity, due for example to other sessions of other MN's. Generally, for whatever reason, the new AP 5 may not be able to provide as much bandwidth as the initial AP 4. These situations could mean the new AP 5 would not be able to fully support the session currently previously handled satisfactorily be the initial AP 4, i.e. it might be that some of the IP data flows of the session might not be supported if handover is implemented.

The new AP 5 performs an admission control algorithm to determine which flows can be accepted, based on the RSVP information received in the Move-Response message. At step s48, the new AP 5 sends a Move-Accept message to the initial AP 4. The Move-Accept message includes the results of the admission control algorithm. At step s50, the initial AP 4 sends a Change-Command message to the MN 6 on its initial channel 14. The Change-Command message also includes the results of the admission control algorithm.

Depending on the admission control algorithm, the results can be either in the form of a simple set of flows that are accepted, implying that they will receive the same QoS level as in the old AP, or may as another possibility also contain a new set of (reduced) QoS parameters per flow, in case the new AP 5 does not have enough resources to maintain the same level of QoS support.

In this example the admission control algorithm results are acceptable to the MN 6. Hence, at step s52, the MN 6 sends an acceptance Change-Response message to the initial AP 4 on its initial channel 14. The MN 6 then switches channels to its new channel 15, awaiting communication via the new AP 5. In other words, the MN 6 has now completed its part of the handover process, and is now waiting for the new AP to start communicating with it. This constitutes the end of the third time period 33. In terms of the data content transmission and reception as represented in FIG. 3 by steps s35 and s51, it will be appreciated that step s51 represents the final data transfer on the initial channel 14.

(In the situation where the admission control algorithm results are not acceptable to the MN 6, then the MN 6 does not send the acceptance Change-Response to the initial AP, and the handover process is terminated. One possibility is to specify that in this event the MN 6 should then try to handover to another AP instead of the AP 5.)

At step s54, the initial AP 4 sends a Move-Command message to the new AP 5, which is in effect a confirmation to the new AP 5 that handover is going ahead. At step s56, in response to the Move-Command message, the new AP 5 sends a Layer-2 (L2) Update Frame to the IP network 2. Steps s54 and s56 take place during a fourth time period 34 as shown in FIG. 3. This fourth time period 34 is characterised as a time period in which the MN 6 has once again switched to its new channel 15 but is not yet in the process of passing data content on that channel.

In response to the Layer-2 (L2) Update Frame, the IP network 2 changes the data path to being switched via the new AP 5 instead of the initial AP 4. This is represented in FIG. 3 by a step s58 in which the content data is passed between the MN6 operating on the new channel 15, the new AP 5, the IP network 2 and the data route 8. In more detail, step s58 comprises: a step s58a in which content data is passed between the MN 6 operating on its new channel 15 and the new AP 5; a step s58b in which content data is passed between the new AP 5 and the IP network 2; and a step s58c in which content data is passed between the IP network 2 and the data route 8. As before, although presented as a discrete step s58 for consistency in FIG. 3, it will again be understood that this represents an

ongoing two-way data flow process during a fifth time period 35, as shown in FIG. 3. This fifth time period 35 is characterised as a time period in which the overall handover process has been completed and content data transfer is taking place via the new AP 5 with the MN 6 operating on its new channel 15.

Broadly speaking, this embodiment tends to provide an advantageous reduction in handover latency compared to the prior art method. This is achieved by virtue of the MN 6 reverting to the initial AP 4 on the initial channel 14 after having sent the Reassociate-Request message to the new AP 5 on the new channel 15 at step s34 during the second time period 32, thereby allowing content data flow to take place during the third time period 33 while the new AP 5, the initial AP 4 and the authentication unit 3 perform handover steps s36-s50.

This can further be appreciated by the following summary of the first to fifth time periods 31-35, as shown in FIG. 3, from the perspective of the MN 6.

In the first time period 31, the MN 6 is communicating with the initial AP 4 on its initial channel 14, including the flow of content data.

In the second time period 32, the MN 6 is switched to its new channel 15 and is communicating a handover request, but not content data, to the new AP 5.

In the third time period 33, the MN 6 is once again communicating with the initial AP 4 on its initial channel 14, including the flow of content data.

In the fourth time period 34, the MN 6 has once again switched to the new channel 15 and is awaiting further communication with the new AP 5, i.e. no content data flows in this fourth time period 34.

In the fifth time period 35, the MN 6 is communicating with the new AP 4 on its new channel 15, including the flow of content data.

Thus, it will be appreciated that content data flow takes place in the first time period 31, the third time period 33, and the fifth time period 35, but not in the second time period 32 and the fourth time period 34. This means there are two time periods when content data flow does not take place, whereas in the prior art method of FIG. 2 there is just one time period when content data flow does not take place, namely second time period 22 of FIG. 2. However, the second time period 32 and the fourth time period 34 of the embodiment of FIG. 3 are singly and in total typically much shorter than the second time period 22 of the prior art example of FIG. 2, thus typically providing an advantageous reduction in the total time that content data flow does not take place compared to the prior art approach. This tends to allow for transmission disruption to be alleviated, which is particularly useful when the content data comprises multimedia data.

The above described embodiment includes use of an RSVP context and an admission control algorithm. However, these aspects are optional, and in other embodiments are not employed. In such an embodiment which is otherwise the same as the embodiment described above, this results in the following differences compared to the above embodiment described with reference to FIG. 3:

(i) at step s46, there is no inclusion of an RSVP context in the Move-Response message;

(ii) the new AP 5 does not perform an admission control algorithm;

(iii) at step s48, there is therefore no inclusion of admission control algorithm results in the Move-accept message;

(iv) likewise, at step s50, there is therefore no inclusion of admission control algorithm results in the Change-Command message;

(v) step s52, in which the MN 6 sends an acceptance Change-Response message to the initial AP 4, will take place by default, or at least it will not be dependent upon an admission control algorithm result.

In a further modified form of any of the above described embodiments, the need for content data flow to be interrupted during the fourth time period 34 (i.e. while the Move-Command message and L2 Update Frame are sent at steps s54 and s56 respectively) may be reduced or removed as follows. In the above described embodiment, the MN 6, having sent the acceptance Change-response message on its initial channel 14 at step s52, then switches channel to the new channel and awaits completion of steps s54 and s56. In the modified version, the MN 6, having sent the acceptance Change-response message on its initial channel

14 at step s52, then remains on the initial channel 14 for a further amount of time so that content data flow may continue on the initial channel for some or all of the time it takes the other system elements to complete steps s54 and s56. This further amount of time may be a predetermined amount of time, or may be determined according to an algorithm using knowledge of previous times taken for steps s54 and s56 to be completed. Another possibility is that the MN6 only remains on the initial channel 14 for the further amount of time if the signal quality is sufficient, i.e. above a predetermined threshold.

In the above described embodiment, the example of handover from AP 4 is given, where the AP 4 is termed the initial AP 4 as the example is described of a scenario where access is initially provided via the AP 4. It will however be appreciated that the handover process described in the above embodiment may also be applied, and forms part of the invention, when applied to an AP connection that is not necessarily the initial, i.e. first AP connection of a particular data session. In other words, for example, the above described process may equally be applied to a following handover from the new AP 5 to a further AP, not shown in FIG. 1, and so on. As such, it is to be understood that the terminology "initial AP", and likewise the terminology "initial channel", is used herein, including in the appended claims, for convenience to identify the initial AP or channel so far as the current handover process is concerned, and is not restricted to meaning the initial AP or channel so far as any other aspect, such as the current data session or call as a whole, is concerned.

In the above described embodiments, the invention is applied to a WLAN operating according to an adapted version of the IEEE 802.11 standard, and the handover process is a modified form of the Inter Access Point Protocol (IAPP). However, in other embodiments the invention may be applied to WLANs (or other access networks, see the following paragraph) operating according to any other suitable standards or protocols, including handover protocols. In this respect, it is noted that the invention is particularly advantageous when implemented in networks where so-called soft handover is not available (such as the IEEE 802.11 standard WLAN in the above embodiment). Nevertheless, the invention is still applicable to and may still be employed advantageously in other types of access network even when soft handover is in fact available. In these situations, the present invention may for example be employed when the soft handover process is either unavailable or of diminished performance, for whatever reason, or in any other form of combination or substitution with soft handover.

In the above described embodiments, the network in which the invention is applied is a WLAN. In yet further embodiments, the invention is applied to a wireless wide area network (WWAN) instead of, or in addition to, a WLAN. Generally, WLAN's and WWAN's may be considered as two types of access networks, and the invention may be applied to any access network.

In the above described embodiments, the WLAN comprises an IP network. It will be appreciated that the IP network is just one example of types of transport networks (for example in an office block, or various office locations of

an organisation, or e.g. throughout an airport, and so on) to which the present invention may be applied. For example, in other embodiments the transport network may be a specific form of network, e.g. a communications network operating according to IP, for example allowing voice over IP communication.